
This document is available in PDF and HTML

<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#61>

PDF, generated from the HTML version, may be out of date.

A shorter version originally appeared in *Computing in Schools*, 1981, and was later extended and published in 1984 as

'Experiencing Computation: A Tribute to Max Clowes,' in *New horizons in educational computing*, Ellis Horwood Series In Artificial Intelligence, Ed. Masoud Yazdani, pp. 207--219, Ellis Horwood, Chichester,

<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#61>

This tribute was much enlarged in March and April 2014, including

-- recollections by Wendy Manktellow (nee Taylor) [\[WMI\]](#), and

-- a draft annotated biography of Max Clowes with publications [\[BIO\]](#).

(Please send corrections and additions to [a.sloman\[at\]cs.bham.ac.uk](mailto:a.sloman[at]cs.bham.ac.uk))

Experiencing computation

A tribute to Max Clowes

by

Aaron Sloman [\[AS\]](#)

<http://www.cs.bham.ac.uk/~axs>

Previously: Reader in Philosophy and Artificial Intelligence

Cognitive Studies Programme [\[COGS\]](#)

School of Social Sciences,

The University of Sussex

[\[Contents\]](#)

NOTES:

Since this was posted online at Birmingham University in 2001, there have been a number of modifications. If you have information about Max's biography that you would be willing to contribute, please let me (AS) know.

Major extensions:

March 2014

Wendy Manktellow worked with Max at Sussex University, helping with project administration. She was then Wendy Taylor, and will be remembered by several of Max's students and collaborators. In February 2014 she stumbled across this web page, and wrote to me with an anecdote which I've appended below [\[WMI\]](#), with her permission.

Added April 2014

[\[BIO\]](#). Draft annotated biography/bibliography of Max Clowes, with help from colleagues



CONTENTS

- [Introduction](#)
 - [Choosing a language](#)
 - [More than just a language](#)
 - [Computing in Schools?](#)
 - [Acknowledgments](#)
 - [Original References](#)
 - [Footnotes to original tribute \(2001\)](#)

 - [APPENDIX 1:
Wendy Taylor/Manktellow
Max's first pupil \(added 9 Mar 2014\)](#)

 - [APPENDIX 2: Draft Incomplete Annotated Biography/Bibliography:
Life and work of Max B Clowes
\(added 11 Apr 2014\)](#)
 - [About this document ...](#)
-

Experiencing Computation: A Tribute to Max Clowes

Originally published in 1984, in *New horizons in educational computing*, Ellis Horwood Series In Artificial Intelligence, Ed. Masoud Yazdani, pp. 207--219, Ellis Horwood, Chichester, A shorter version originally appeared in *Computing in Schools*, 1981.

Introduction

Max Clowes died of a heart attack on Tuesday 28th April 1981. He was one of the best known British researchers in Artificial Intelligence, having done pioneering work on the interpretation of pictures by computers. His most approachable publication is cited below. He was an inspiring teacher and colleague, and will be remembered for many years by all who worked with him. He helped to found AISB, the British society for the study of Artificial Intelligence and the Simulation of Behaviour, now expanded to a European society. This tribute is concerned mainly with his contribution to education.

He was one of the founder members of the Cognitive Studies Programme begun in 1974 at the University of Sussex, a novel attempt to bring together a variety of approaches to the study of Mind, namely Psychology, Linguistics, Philosophy and Artificial Intelligence. During the last few years his interests centred mainly on the process of teaching computing to absolute beginners, including those without a mathematical or scientific background. He was one of the main architects of the Sussex University POP11 teaching system (along with Steve Hardy and myself), which has gradually evolved since 1975. In this brief tribute, I shall sketch some main features of the system, and hint at the unique flavour contributed by Max.

POP11 embodies a philosophy of computer education which is relatively unusual. It includes a language, a program-development environment, and a collection of teaching materials including help facilities, much on-line documentation, and a large collection of exercises and mini-projects. Unfortunately, it is at present available only on a PDP11 computer running the Unix operating system, though a version now being written in C should be available for use on a VAX by the end of this year.[\[CPOP\]](#)

When we started planning the system, in 1974, we were much influenced by the writings of John Holt (see references at end), the work on LOGO at MIT by Seymour Papert and colleagues, and at Edinburgh University by Sylvia Weir, Tim O'Shea, and Jim Howe.

These influenced our conviction that learners of all ages should be treated not like pigeons being trained by a schedule of punishment and reward, but like creative scientists driven by deep curiosity and using very powerful cognitive resources. This entailed that learners should not be forced down predetermined channels, but rather provided with a rich and highly structured environment, with plenty of opportunities to choose their own goals, assess their achievements, and learn how to do better next time through analysis of failures.

Although these needs can to some extent be met by many older learning environments (e.g. meccano sets, learning a musical instrument, projects), the computer seemed to be potentially far more powerful, on account of its speed, flexibility, reactivity and ability to model mental processes. Instead of making toy cranes or toy aeroplanes, or dolls, students could make toy minds.

Unlike educational philosophies which stress 'free expression', this approach stresses disciplined, goal oriented, technically sophisticated, activities with high standards of rigour: programs will not work if they are badly designed. Yet the computer allows free expression to the extent that students can choose their own goals, and their own solutions to the problems, and the computer will patiently, more patiently than any teacher, pay detailed attention to what the student does, and comment accordingly, by producing error messages, or running the program and producing whatever output is required. Of course, error messages need to be far more helpful than in most programming environments, and the system should make it easy for the student to make changes, to explore 'where the program has got to', and to try out modifications and extensions without a very lengthy and tedious edit compile and run cycle.

These ideas are embodied in a course, *Computers and Thought*, offered as an unassessed optional first year course for students majoring in Humanities and Social Science subjects. By making the computer do some of the things people can do, like play games, make plans, analyse sentences, interpret pictures, the students learn to

think in a new way about their own mental processes. Max put this by saying that he aimed to get students to 'experience computation' and thereby to 'experience themselves as computation'.[\[CT\]](#) In other words, our answer to the student's question 'Does that mean I'm a computer?' is 'Yes'. Of course people are far more intricate, varied, flexible, and powerful than any man-made computer. Yet no other currently available framework of concepts is powerful enough to enable us to understand memory, perception, learning, creativity and emotions.

Choosing a language

Part of the LOGO philosophy was that beginners playing with computers needed a very powerful language, making it easy to perform interesting tasks quickly. (See Papert). Interesting tasks include building 'toy minds'. Thus we needed a language which is interactive and provides procedures with recursion and local variables, and facilities for non-numerical problem solving, such as list manipulation. This rules out BASIC.

LOGO is much more powerful, but, we felt, did not go far enough: after all, it was designed for children, so it could not be powerful enough for children! PASCAL was ruled out as too unfriendly and even less powerful than LOGO. (E.g. the type structure makes it impossible to program a general-purpose list-processing package: the package has to be re-implemented for numbers, words, lists etc.) We did not know about APL. It might have been a candidate, though its excessively compressed syntax which delights mathematicians is hardly conducive to easy learning by the mathematically immature. Moreover it appears to have been designed primarily for mathematical applications, and does not seem to have suitable constructs and facilities for our purposes. PROLOG would have been considered had an implementation been available, though it too is geared too much towards a particular class of problems, and is hard to use for others.

We therefore reduced the choice to LISP and POP2, and settled for the latter because of its cleaner, more general semantics (e.g. functions are ordinary values of variables), more natural syntax, and convenient higher-level facilities such as partial-application and powerful list-constructors (Burstall et. al). A subset of POP2 was implemented on a PDP11/40 by Steve Hardy, and then extended to provide a pattern-matcher, database, and other useful features. We now feel that the power of PROLOG (see Kowalski 1979) should be available as a subsystem within POP, and are planning extensions for the VAX version.[\[CM\]](#)

More than just a language

From the start we intended to provide a wide range of facilities in the library, so that students could easily write programs to do things which interested them: draw pictures, analyse pictures, play games, have conversations relating to a database of

knowledge, etc. We soon also found the need for help-facilities, on-line documentation of many kinds, and a simple, non authoritarian teaching program (written in POP11, and calling the compiler as a subroutine to execute the student's instructions), which could, for some learners, feel less daunting than a twenty page printed handout.

One of the ideas that played an important role in our teaching strategy was an analogy between learning a programming language, and learning a natural language, like English. The latter does not require formal instruction in the syntax and semantics of the language. The human mind seems to possess very powerful capacities for absorbing even a very complex formalism through frequent and fruitful *use*. So, instead of starting with lectures on the language, we decided to give the students experience of using the language to get the computer to do things we hoped would make sense to them. So they were encouraged to spend a lot of time at the terminal, trying out commands to draw pictures, generate sentences, create and manipulate lists, etc, and as they developed confidence, to start working towards mini-projects.

Extending or modifying inadequate programs produced by the teacher (or other students) provides a means of gaining fluency without having to build up everything from the most primitive level. Naturally, this did not work for everyone. Some preferred to switch at an early stage to learning from more formal documentation. Some found the pain of even minor errors and failures too traumatic and needed almost to be dragged back to try again - often with some eventual success. Some, apparently, had insuperable intellectual limitations, at least within the time-scales available for them to try learning to program. But many students found it a very valuable mind-stretching experience.

One of the ways in which Max contributed to this was his insistence that we try to select approaches and tasks which were going to be more than just a trivial game for the student. He was able to devise programming exercises which could be presented as powerful metaphors for important human mental processes - such as the pursuit of goals, the construction of plans, the perception and recognition of objects around us and the interpretation of language. He would start the 'Computers and Thought' course by introducing students to a simple puzzle-solving program and erect thereon a highly motivating interpretation: treating it as a microcosm of real life, including the student's own goal-directed activities in trying to create a working program. (Perhaps this is not unconnected with a slogan he defended during his earlier work on human and machine vision: "Perception is controlled hallucination" - hallucinating complex interpretations onto relatively simple programs helps to motivate the students and give them a feel for the long term potential of computing).

Moreover, he always treated teaching and learning as more than just an intellectual process: deep emotions are involved, and need to be acknowledged. So he tried to help students confront their emotions of anxiety, shame, feeling inadequate, etc., and devised ways of presenting material, and running seminars, which were intended to help the students build up *confidence* as well as understanding and skills.

There is no doubt that for many students the result was an unforgettable learning experience. Whether they became expert programmers or not, they were changed persons, with a new view of themselves, and of computing. Some of this was due to his inimitable personality. In addition he advocated strategies not used by many university teachers at any rate: such as helping all the students in a class to get to know one another, prefacing criticisms with very encouraging comments, and helping students cope with their own feelings of inadequacy by seeing that others had similar inadequacies and similar feelings about them, whilst accepting that such 'bugs' were not necessarily any more permanent than bugs in a computer program.

As a teacher I found myself nervously treading several footsteps behind him - too literal-minded to be willing to offer students his metaphors without qualification, yet benefitting in many ways from his suggestions and teaching practice.

Just before he left Sussex we had a farewell party, at which he expressed the hope that we would never turn our courses into mere computer science. There is little risk of that, for we have learnt from him how a different approach can inspire and motivate students. The computer science can come at a later stage - for those who need it. For many, who may be teachers, managers, administrators, etc. rather than programmers or systems analysts, the formalities of computer science are not necessary. What is necessary is a good qualitative understanding of the range of types of things that can be done on computers, and sufficient confidence to face a future in which computation in many forms will play an increasingly important role.

None of this should be taken as a claim that the teaching system based on POP11, used by Max and the rest of us, is anywhere near perfect. We are aware of many flaws, some of which we feel we can remedy. But there is still a great deal of exploring to be done, in the search for a good learning environment, and good learning experiences. Moreover, we don't know how far what is good for novice Arts university students would also be good for school children, though several have played with our system and enjoyed it. We are still in the process of improving the POP virtual machine to make it a more natural tool for thinking about processes. This is a never-ending task. Probably a language is needed which can be 'disguised' to present a simpler interface for the youngest learners, without sacrificing the power to do interesting things very quickly. To some extent the 'macro' facility in POP (see Burstall et. al.) makes this possible.

Computing in Schools? (1980 version)

In December 1980 Max left Sussex, to join a project on computing in schools. Although I don't know exactly what his plans were, I feel that he would probably have fed many important new ideas into the educational system. New ideas are surely needed, for teaching children to program in BASIC is like teaching them to climb mountains with their feet tied together: the permitted steps are so very small. Moving towards COMAL will merely loosen the ropes a little. Teaching them PASCAL will loosen the ropes a lot more, but add heavy weights to their feet: taking some big steps is possible in PASCAL, but unnecessarily difficult. And some things are not possible, as mentioned previously.

The problems of providing better teaching languages and teaching environments are enormous, since available computers are much too small, and there is too little expertise on tap in schools. I think Max might have tried to press for funds to be diverted from stand-alone microcomputers to larger, shared, machines, or networks, making available better languages, shared libraries, larger address spaces and increased opportunities for learners to help one another, including teachers. This may be more expensive: but what could be a more important use of computers?

Such a system, based on a DEC 11/23, running the UNIX operating system, has just been introduced at Marlborough College, partly as a result of Max's inspiration.

It will be interesting to watch what happens there. Maybe we'll learn from that how to use the newer, cheaper, bigger, faster systems that will shortly be on the market. Let us hope the existing puny machines with their inadequate languages will not have turned too many people off computing for life.

[End of original tribute]

Note added 7 Sep 2014

Something like that feared disaster happened in the three decades after those comments were written.

A new UK Computing At School movement was initiated around 2009, <http://computingatschool.org.uk/> and by 2014 the UK school computing educational system has been considerably shaken up. I sometimes wonder whether this might have happened sooner if Max had not died so young.

Acknowledgments (in the original obituary notice)

I am grateful to Reg Baukham and Masoud Yazdani, who reminded me of some features of Max's teaching omitted from a first draft, and who pointed me at things he had written in CSGAS, an informal magazine edited by Masoud Yazdani.

Original References (1981)

R M Burstall, et. al. *Programming in POP2*, Edinburgh University Press, 1972

Max Clowes, 'Man the creative machine: A perspective from Artificial Intelligence research', in J. Benthall (ed) *The Limits of Human Nature*, Allen Lane, London, 1973.

John Holt, *How Children Learn*, Penguin Books

John Holt, *How Children Fail*, Penguin Books.

Robert Kowalski *Logic for Problem Solving*, North Holland, 1979.

Seymour Papert, *Mindstorms*, Harvester Press, 1981.

<http://hopl.murdoch.edu.au/showlanguage2.prx?exp=7352>

Entry for Max Clowes at HOPL (History of Programming Languages) web site. Alas, it now appears to be defunct. (12 Apr 2014) Try this instead:

<http://archive.today/hopl.murdoch.edu.au>

FOOTNOTES TO ORIGINAL TRIBUTE

(Added 2001, for Web version)

[AS]

Now at the University of Birmingham <http://www.cs.bham.ac.uk/~axs/>

I discovered a version of this article among some old files in 2001, and thought it would be useful to make it available online, at least for people who remember Max, and perhaps others.

I added a few footnotes, putting the text in the context of subsequent developments.

Note added 28 Mar 2014 I have been working on a paper on unsolved problems in vision, making use of some of Max's ideas, here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vision>

[COGS]

The Cognitive Studies Programme, was founded by Max Clowes and others at the University of Sussex as a cross-disciplinary undergraduate programme within the School of Social Studies. Later, a few years after he died it became a separate School of Cognitive Sciences. Within a few years it had enlarged further and become what is now known as COGS, The School of Cognitive and Computing Sciences, with an international reputation. But for the early inspiration and influence of Max Clowes it would never have existed. My own role in its development was a direct consequence of the influence Max had in changing my research direction, from about 1969. Others who played major roles in the formation and growth of COGS were also inspired by Max, sometimes indirectly.

[CPOP]

Note added 2001: In fact the version in C proved to be a temporary bridge towards a Pop11 in Pop11 compiler subsequently developed by John Gibson, which became the core of the multi-language Poplog system, which was later sold world wide. It is now available free of charge with full system sources at <http://www.cs.bham.ac.uk/research/poplog/freepoplog.html> where much of the still useful AI teaching material has its roots in his approach to teaching. E.g. <http://www.cs.bham.ac.uk/research/poplog/teach/river>

[CT]

Some years later, ideas that had developed out of the course appeared in a book: Sharples, M. Hogg, D. Hutchison, C. Torrance, S. Young, D. *Computers and Thought: a practical introduction to Artificial Intelligence* MIT Press 1989. Many students without backgrounds in programming and mathematics find that this book provides a useful route into AI and computational cognitive science. The text (without images) and Pop-11 code can be downloaded here: <http://www.cs.bham.ac.uk/research/projects/poplog/computers-and-thought/>

[CM]

Chris Mellish implemented a Prolog in Pop-11 in 1981, and after that the system combining both languages came to be known as Poplog. Later, other languages were added, including Common Lisp, and Standard ML. There are Wikipedia entries with more information:

-- <http://en.wikipedia.org/wiki/Poplog>

-- <http://en.wikipedia.org/wiki/POP-11>

[WM] APPENDIX 1:

Wendy Taylor/Manktellow

Max's "first pupil"

(added 9 Mar 2014)

In February 2014, out of the blue, I received a message from Wendy Manktellow, who, as Wendy Taylor, had worked with Max as an administrative assistant while he was at Sussex University. She described an episode that sheds light on Max as teacher and colleague, which I have added here, with her permission.

START LETTER

I have just found your tribute to Max on line and realised, as I read about his thinking on teaching and learning, that I was his first pupil!!!

I recall his growing consternation having been awarded his chair, that he was to teach Arts students who would have little or no maths. So he decided to use me as his guinea pig.

There was one session when he was trying to teach me how to display a loop on screen from digitised cursive-script handwriting (presumably received from Essex who were working on that area of AI at the time). We were both getting very stressed as I displayed it every way but upright. Eventually, in tears, I said: "Max I am just too stupid to learn."

He stared at me for a moment and said: "No Wendy, I am too stupid to teach."

Years later, after my English degree and a PGCE (both Sussex) when I was teaching at Comprehensive level, that incident kept coming into my mind. It was a eureka moment about the nature of teaching and learning.

That day, when I'd dried my tears, Max and I had a long talk about it, about the feelings of inadequacy, shame and stupidity in the teaching and learning process. I think we were both stunned that we were both feeling the same emotions.

Then we got back to finding loops and we both got it right!!!

I never forgot what I learned that day and it made such a huge difference throughout my teaching career.

Yours is a fine tribute. Max was such a very special person. You all were. Those pioneering days of the Cognitive Studies Programme in the prefabs next to the refectory were so full of energy and excitement. I have wonderful memories of those years.

Thank you for finding the time to read this.

Wendy Manktellow (nee Taylor)

END LETTER

[BIO] APPENDIX 2:

Draft Incomplete Annotated Biography/Bibliography

Added: 11 Apr 2014;

Updated: 11 Mar 2016 (formatting)

Feb 2016: Added more quotes, about metaphor and Gombrich in section on [ICA article](#).

31 May 2015: slightly modified [Reutersvard figure](#).

12 Sep 2014 (Added Reutersvard cubes example.)

16 Apr 2014; 5 May 2014; 7 Sep 2014; 9 Sep 2014 (Added Boden's review);

With thanks to Margaret Boden, Steve Draper, Mark Jenkin, Alan Mackworth, Keith Oatley, Frank O'Gorman, Vaughan Pratt, Mark Steedman, Tom Vickers, and google.

Life and work of Max B Clowes

Contents of Annotated Biography/bibliography:

- [\[Back to Obituary Notice \(above\)\]](#)
- [Biography/Bibliography \(with comments\)](#)
- [Date of Birth 1933](#)
- [PhD in the department of Physics, University of Reading.](#)
- [M.B. Clowes & R.W. Ditchburn \(1959\)](#)
- [NPL Teddington 1959 - 1963](#)

- [Clowes, M. B., & Parks, J. R.\(1961\).](#)
 - [Oxford Psycho-linguistics Research Unit 1963-196?](#)
 - [Founding AISB \(1964\)](#)
 - [M. B. Clowes, Perception, picture processing and computers](#)
 - [Biology and Engineering](#)
 - [Computing Research Section, C.S.I.R.O., Canberra, Australia. 19??-1969?](#)
 - [M.B. Clowes, Pictorial relationships - a syntactic approach.](#)
 - [Influence of Chomsky](#)
 - [Seeing as hallucinating](#)
 - [Images vs things depicted](#)
 - [Influences and connections while at CSIRO, Australia](#)
 - [Move to Sussex University, Experimental Psychology, 1969](#)
 - [First PhD student, Robin Stanton](#)
 - [Margaret Boden's review](#)
 - [On Seeing Things \(1971\)](#)
 - [Controlled hallucination?](#)
 - [IJCAI 1971](#)
 - [M.B. Clowes, "Scene analysis and picture grammars", \(1972\)](#)
 - [M.B. Clowes, ICA Lecture and book chapter 'Man the creative machine: A perspective from Artificial Intelligence research'\(1972/3\)](#)
 - [O'Gorman and Clowes Collinearity \(1973, 1976\)](#)
 - [Other publications to be added:](#)
 - [Students and colleagues at Sussex](#)
 - [Obituary notices](#)
-

Biography/Bibliography (with comments)

- **Date of Birth 1933**
(According to Margaret Boden *Mind as Machine*, p.787)
- **PhD in the department of Physics, University of Reading**, supervised by R.W. Ditchburn (1958 or 1959?)
- The following publication seems to have been based on the PhD thesis:

M.B. Clowes & R.W. Ditchburn (1959)

An Improved Apparatus for Producing a Stabilized Retinal Image, *Optica Acta: International Journal of Optics*, 6:3, pp 252-265, Taylor & Francis, DOI: 10.1080/713826291 <http://dx.doi.org/10.1080/713826291>

Abstract:

Criteria for defining the efficiency of an apparatus for stabilizing the retinal image are formulated. A distinction is made between geometrical stabilization and stabilization of illumination. A new technique is described which employs a telescopic normal incidence system. This makes it possible to obtain geometrical compensation both for rotations and for translations of the eye. It also gives good illumination stabilization. The degree of compensation achieved may be evaluated by precise physical measurements. About 99.7 per cent of natural eye rotations in horizontal and vertical planes is compensated and the effect of translations is negligible. The apparatus is designed to permit easy interchange of normal

and stabilized viewing conditions.

NOTES:

(a) The Acknowledgments section states:

"We would like to thank Dr. D. H. Fender and Dr. Stella Mayne for much helpful discussion. Mr. W. S. Martin has provided technical assistance in the construction of the apparatus. This work was supported by a research grant No. B-1233 from the Department of Health, Education and Welfare, Public Health Service, U.S.A."

(b) The document states for Max:

"Now at the National Physical Laboratory, Teddington."

- **NPL Teddington 1959 - 1963**

(Dates and NPL information provided by Tom Vickers.)

He seems to have gone from Reading to the National Physical Laboratory, Teddington (NPL), where he worked with John Parks, and David Yates.

I don't have access to the next two papers, apparently written at NPL and referenced in his 1967 paper, below.

- **Clowes, M. B., & Parks, J. R.(1961)** A new technique in character recognition. *Computing Journal* , 4, 121-126.
- Clowes, M.B.(1962). The use of multiple auto-correlation in character recognition. *Optical Character Recognition*. Fischer, Pollack, Radack & Stevens, Eds. Pp. 305-318. Baltimore: Spartan Books, Inc.

- **Oxford Psycho-linguistics Research Unit 1963-196?**

From NPL he moved to the MRC Psycho-linguistics Research Unit at Oxford. The 1967 book chapter was written there.

- **Founding AISB (1964)**

According to Margaret Boden in *Mind as Machine* Vol 1, p.364, in 1964 Max Clowes, presumably still in Oxford, started a "Study Group on Artificial Intelligence and the Simulation of Behaviour (AISB)" as a subgroup of the British Computer Society, and writes "within a decade this had become the AISB Society, under the interdisciplinary leadership of researchers at the universities of Edinburgh, Sussex, and Essex. The name was a psychological pun on Clowes' part: 'A is B'".

AISB is still flourishing <http://www.aisb.org.uk/> and celebrated its 50th year in 2014: www.aisb.org.uk/events/aisb14 .

Note: AI did not start at Sussex until Max arrived around 1969, and it did not start at Essex until Pat Hayes and Mike Brady, and later Yorick Wilks, established it around 1972 and after. So initially the "leadership" must have involved Edinburgh plus Max Clowes, then at Oxford? PDF versions of the conference proceedings (from 1974) are available at <http://www.aisb.org.uk/asibpublications/convention-proceedings>

- **M. B. Clowes, Perception, picture processing and computers** *Machine Intelligence Vol 1* pp 181--197 Eds. N L Collins and Donald Michie, Oliver & Boyd, 1967. <http://aitopics.org/sites/default/files/classic/Machine%20Intelligence%201%262/MI1%262-Ch.12-Clowes.pdf> (Later re-published by Edinburgh University Press.)

For information about the Machine Intelligence series see
<http://www.doc.ic.ac.uk/~shm/MI/mi.html>

Max's address is given as M.R.C. Psycho-Linguistics Research Unit, University of Oxford
The ACKNOWLEDGMENTS section states:

"I would like to acknowledge with gratitude the use of the computing facilities at the Culham Fusion Research Laboratory, U.K.A.E.A., the assistance of Douglas Brand and Barry Astbury, and valuable discussions with John Marshall, Professor N. S. Sutherland and Professor R. C. Oldfield."

A footnote says:

Present address: Computing Research Section, C.S.I.R.O., Canberra, Australia.

So presumably he went from Oxford to Australia. When?)

- **Biology and Engineering**

Although the paper was written in an engineering context, Max's concerns, unusually for AI researchers at that time, were partly biological, as illustrated by these extracts:

"The ability to interpret and respond to the significant content of pictures is shared by a large proportion of the animal kingdom and a small but growing number of machines."

....

"No machine can as yet approach this sort of performance. We may well ask, therefore, what if anything we can learn from a study of 'machine perception' which could conceivably help us to understand human perception. The answer lies in the fact that any realistic account of the mechanism underlying perceptual or any other human skills will be complex. The virtue of a computer lies in its ability to capture in a definite form processes of indefinite complexity and subtlety, and moreover, to permit an evaluation of the efficacy of the proposed description by trying it out on the actual task."

He discusses the work done in linguistics on formally characterising linguistic structures (e.g. spoken or written sentences) at different levels of abstraction, and remarks:

"Comparable mechanisms for processing pictures have not appeared to any significant extent either in Artificial Intelligence or in Psychology. Instead the emphasis has been on classifying pictures, i.e., upon the equivalent of deciding the type (S1 or S2) of a sentence. The possibility that a structural description of a pictorial object is necessary has only recently emerged in Artificial Intelligence (Kirsch 1964). It appears to be crucial to the automation of many picture processing tasks (e.g., interpreting bubble-chamber photographs, recognising fingerprints). That perceptual behaviour involves description as well as classification could be supported by innumerable examples beyond the two already quoted in the introduction to this paper. However, it seems to have received scant attention in recent psychological literature. This may well be because overtly descriptive (rather than classificatory) behaviour has a large verbal element. There is a strong temptation to avoid this 'complication' by designing experiments which merely require simple YES/NO responses"

(... or selecting labels from a fixed set, he might have added).

Max's comment remains relevant to a great deal of 21st Century AI research in machine vision (i.e. up to 2014 at least), focusing on training machines to attach labels as opposed to understanding structures (I would now add "and processes involving interacting structures").

One of the points he could have made but nowhere seems to have made, is that natural vision systems are mostly concerned with motion and change, including change of shape, and change of viewpoint. The emphasis on static scenes and images may therefore conceal major problems e.g. those pointed out by J.J.Gibson. AI vision researchers later started to address this (partly influenced by Gibson), though as far as I can tell, it never interested Max.

So he specifies the objective of

"... development of a formal (i.e., rigorous) theory of picture processing which aims to provide a structural description of pictorial objects as well as a classification of them (wherever appropriate)."

- **Computing Research Section, C.S.I.R.O., Canberra, Australia. 19??-1969?**
(Commonwealth Scientific and Industrial Research Organisation www.csiro.au/)
The next publication I've found gives Max's address as "Division of Computing Research, C.S.I.R.O., Canberra"
- **M.B. Clowes, Pictorial relationships - a syntactic approach.** *Machine Intelligence Vol 4*, pp 361-384, 1969 B. Meltzer & D. Michie (eds.). Edinburgh University Press
<http://aitopics.org/sites/default/files/classic/Machine%20Intelligence%204/MI4-Ch20-Clowes.pdf>
This paper introduced some important new ideas to the study of vision, partly under the influence of Noam Chomsky.
- **Influence of Chomsky**
This important paper was deeply influenced by Chomsky, especially Chomsky,N.(1965) *Aspects of the theory of syntax*. Cambridge, Mass: MIT Press.

In particular Max (like several AI vision researchers in that decade) argued that images had a type of syntax and what they depicted could be regarded as semantic content. Max attempted to develop a research methodology inspired partly by Chomsky's work, emphasising the importance of concepts of

- 'ambiguity' (two possible semantic interpretations for one syntactic form)
- 'paraphrase' (two syntactic forms with the same semantic content),
- 'anomaly' (syntactically well formed images depicting impossible semantic contents -- impossible objects, e.g. ["The devil's pitchfork"](#) often misdescribed as an "illusion"!)

He also emphasised important differences between pictures and human languages, e.g.

"The variety of relationships which we can readily identify and name is much greater in pictorial expressions than in string expressions."

And goes on

"Published accounts of 'picture syntax' have not provided any systematic accounts of the variety of pictorial relationships with which they deal, much less a discovery procedure for those relationships. This omission may of course be intentional in the sense that no attempt is being made to capture our intuitive knowledge of picture structure in these picture syntaxes. In this account, however, we adopt as goal the formal description of our pictorial intuitions and accordingly we shall adopt a more or less systematic methodology for ascertaining what these intuitions are."

In section 4.5 he writes *"we are characterising our intuitions about picture structure, not erecting some arbitrary picture calculus."* This leads to the notion that the same picture, or portion of a picture, may instantiate different qualitative, relational, structures at the same time, i.e. different 'views'.

"Significantly, however, we cannot hold these multiple views simultaneously -- we switch between them. Formally, that is, we can only assign a single relational structure at a time, although this structure may relate a number of items ... in quite a complex manner."

Unlike the majority(?) of current computer vision researchers he did not simply accept the properties and relationships that are derivable via standard mathematical techniques from image sensor values and their 2-D array co-ordinates (e.g. defining "straightness" in terms of relationships between coordinates in a digitised image) but instead attempted to identify the properties and relationships that are *perceived* (consciously or unconsciously) by human viewers and *used* in interpreting visual contents, i.e. working out what is depicted (the semantics).

This approach has important consequences:

"When faced, however, with a wholly novel picture, for example that produced in some esoteric experiment in physics, we may find that it takes some considerable time to adjust our view so as to recover the significant elements of Form and reject the insignificant."

...

"The conclusion we would draw from this is that the structure we assign to a picture is determined not solely by the 'raw data' of that picture but also by a priori decisions as to the varieties of relationship we expect to find there. The question therefore becomes 'can we formalise these a priori decisions?'"

There is far more in that important paper, including discussions of how perceived space might be structured in ways that depend on relationships within the percepts rather than some absolute coordinate system, often assumed by researchers on machine vision, then and now. His ideas here seem to be the basis of his later claims that perception involves creativity and metaphor in his 1971 presentation at the ICA ([below](#)). (Compare Boden's comments, in her brief review mentioned [below](#).)

- **Seeing as hallucinating**

The 1969 paper includes a precursor of the slogan "Perception is controlled hallucination" attributed to him by later authors, discussed [below](#):

"We may summarise the foregoing argument as 'People see what they expect to see'. The essential rider is that what they want to see is things not pictorial relationships, that is, the a priori decisions reflect assumptions about the things and events which we expect to see exhibited in the picture. We shall argue that it is necessary and indeed possible to give a structural characterisation of things and events which is a mapping of the relational structure of the picture. This characterisation we call the semantics of the picture."

There is some ambiguity here as to what sort of contrast is implied by "things not pictorial relationships", compounded by the phrase "exhibited in the picture".

In view of his more explicit discussions in later papers, I take him to be saying here that we expect to see things that are not in the picture but are represented by the contents of the picture. That would include, for example, seeing 3-D structures or object-fragments represented in a picture: the plane surface in which the picture lies can include only 2-D entities and their 2-D relationships.

The 1971 paper ([listed below](#)) is unambiguous on this point: the entities represented in the picture (the picture's "semantic" content) have 3-D structure, namely polyhedra, whose surfaces lie in different planes, most of which are not parallel to the picture plane.

What sorts of entities a collection of lines is intended to denote can affect how it should be parsed. E.g. he points out that in a circuit diagram, straightness of lines, and the existence of corners are less important than they might be in other pictures (e.g. a drawing of a building).

"A machine (or program) capable of mediating translations between these various languages would utilise the underlying semantic structure as the 'pivot' of the translational process. We could describe such a machine as 'informed' -- 'informed', that is, about the varieties of relationship applicable in these various representations of an event. It would not, however, be intelligent. Such an appellation should be reserved for a machine (like us) capable of formulating and testing hypotheses about new relationships and ultimately about new systems of attainable concepts manifesting these relationships."

- **Images vs things depicted**

These ideas about interpreting image structures as depicting non-pictorial structures (e.g. 3-D objects, electrical circuits, and other things that constitute the semantic content of the pictures) are ignored by many current vision researchers (in 2014), who, instead, use powerful mathematical machinery and elaborate machine statistical learning regimes, whose result is a collection of narrowly focused (shallow?) systems that perform well in some (often arbitrary) benchmark tests for recognising or labelling image fragments (e.g. hand-written characters), but are far behind the visual capabilities of a human toddler, a squirrel, a nest-building bird, an elephant, and many other animals, whose evolution and epigenesis are subject to pressures and constraints that are quite different from typical AI vision benchmarks.

I am not sure whether Max drew the conclusion that instead of totally general purpose learning mechanisms applied to the raw data of visual and other sensors, human-like intelligent machines would need to have learning mechanisms tailored to the kinds of environments in which we evolved, and preferences for types of "syntactic" and "semantic" ontologies that have been found useful in our evolutionary history. Research on learning using artificial neural nets may be thought to meet that requirement, but that could be based on misunderstandings of functions and mechanisms of biological brains. Compare John McCarthy on "[The well designed child](#)".

At that time, I don't think Max knew how much he was echoing the viewpoint of Immanuel Kant in *The Critique of Pure Reason* (1781). However, he was aware of the work of von Helmholtz (perception is "unconscious inference") and he may have been aware of M.L.J Abercrombie's influential little book [Abercrombie \(1960\)](#), which made several similar points from the viewpoint of someone teaching trainee zoologists and doctors to see

unfamiliar structures, e.g. physiological fragments viewed in a microscope.

Max later acknowledged the connection between his work and Kant's philosophy in Footnote 2 of the 1971 paper 'On Seeing Things' ([listed below](#)).

- **Influences and connections while at CSIRO, Australia**

The Acknowledgments section of the Machine Intelligence 4 paper states:

"The approach to picture interpretation outlined here has emerged from the Verbigraphics Project. It is a pleasure to acknowledge my indebtedness to my colleagues D. J. Langridge and R. B. Stanton. I am grateful to Dr G. N. Lance for encouraging us and supporting us in this work."

The MSc Thesis of Vaughan Pratt, Dated August 1969, University of Sydney, Title: "Translation of English into logical expressions"

<http://boole.stanford.edu/pub/PrattTransEngLogExpns.pdf> acknowledges "Dr Max Clowes, formerly of CSIRO, Canberra, for arousing my interest in transformational approaches to English", and also acknowledges "Associates of Max, including Robin Stanton, Richard Zatorski, Don Langridge and Chris Barter."

- **Move to Sussex University, Experimental Psychology, 1969**

In Oxford, Max had worked with Stuart Sutherland, who later came to Sussex University as head of the Experimental Psychology (EP) laboratory in the School of Biological Sciences (BIOLS). This functioned more or less independently of the social, developmental, and clinical psychology groups in schools within the Arts and Social Sciences "half" of the University.

A result of Sutherland's arrival was that Max was invited to return to the UK to a readership in EP, where he arrived in 1969. Somehow I came to know him and he, Keith Oatley and I had a series of meetings in which we attempted to draft a manifesto for a new multi-disciplinary research paradigm, including AI, psychology, philosophy and linguistics.

- **First PhD student, Robin Stanton**

Robin Stanton, mentioned in the acknowledgments section of the MI4 paper, was supervised by Max for a PhD at ANU. He came to Sussex University after Max moved there in 1969, where I believe he completed the PhD in 1970, and remained for a while doing post-doctoral research with Max before returning to Australia. Some of that information was obtained from: http://cci.anu.edu.au/researchers/view/robin_stanton/ (apparently no longer available: 26 Feb 2016)

Robin's work shifted from AI to more "central" computing science thereafter.

- **Margaret Boden's review**

None of us knew at that stage that [Margaret Boden](#), then a lecturer in Philosophy and Psychology at Sussex University, also had a strong interest in AI, and had been reviewing and commenting on AI publications for several years, in Philosophy journals, as reported in [Boden\(2013\)](#). She had even reviewed several of the *Machine Intelligence* volumes edited by Bernard Meltzer and Donald Michie, and in her review of MI4 (1969) in the *British Journal for the Philosophy of Science* ([Boden, 1970](#)), she singled out Max's 1969 contribution ([discussed above](#)) for extended praise and discussion:

"Clowes' paper is a stimulating discussion of what he calls the 'syntax' of pictorial representation and visual perception. Clowes sees his enterprise as closely parallel to that of Chomsky: both are searching for the deep grammar, our intuitions of which determine our understanding of the expression we can readily describe in terms of its surface grammar. (And both believe that the system of attainable concepts may be limited by a priori features-whether perceptual or linguistic.)"

She also noted the connection with Kant's philosophy:

"Clowes suggests a skeleton metalanguage, and tries to state our intuitions of form and position in terms of it. He also uses it to express certain deeper intuitions which he believes -- on empirical grounds -- to be involved in our visual perception: e.g. the use of implicit axes, the implicit assignment of an interval scale with a small (and relatively fixed) number of units, and so on. These basic relationships which, he says, we expect a priori to find in visual patterns, are perhaps less startling than some of the Kantian categories. But, if Clowes is right, they do determine our perceptual interpretations of form and objects on the basis of visual patterns, and they do set limits to the structure of the system of attainable concepts."

- **On Seeing Things (1971)**

This is Max's best known paper, closely related to the talk he gave at the Institute of Contemporary Arts (ICA) in 1971 or 1972 as part of a series of lectures. A paper based on the talk was published [later](#) in a book based on the lecture series.

M.B. Clowes, On seeing things, *Artificial Intelligence*, 2, 1, 1971, pp. 79--116,
[http://dx.doi.org/10.1016/0004-3702\(71\)90005-1](http://dx.doi.org/10.1016/0004-3702(71)90005-1)

This developed the themes summarised above and introduced the line-labelling scheme used in interpretation of pictures of polyhedra, independently discovered by David Huffman [Huffman \(1971\)](#), and referred to by Clowes in *Footnote 1*, which states:

"Omitted from this discussion is any treatment of Huffman's [24] excellent study. The two papers employ what is essentially the same idea -- the interpretation of junctions as edge complexes -- in a rather different manner. Huffman's treatment is applied to single bodies, not to scenes as in the present paper, but contains a more extensive formulation of the concept of a well-formed polyhedron. The algorithm reported here is not paralleled by any procedural mechanism in Huffman's account."

This shared idea is often referred to as "Huffman-Clowes" labelling, and was generalised by many later researchers, including David Waltz who enriched the ontology and showed that constraint propagation could often eliminate the need for expensive search, and Geoffrey Hinton who showed that use of probabilities and relaxation instead of true/false assignments and rigid constraints, allowed plausible interpretations to be found in the presence of noise (e.g. missing or spurious line fragments or junctions) that would not be found by the alternative more 'rigid' mechanisms.

One of the themes of the paper reiterates the syntax/semantics distinction made in his earlier papers, emphasising the need for different domains to be related by the visual system, e.g. the picture domain and the scene domain, also referred to as the 'expressive' and 'abstract' domains. Consistency requirements in the scene (abstract) domain constrain the interpretation of the previously found structures in the picture (expressive) domain. An example in the paper is that in a polyhedral scene an edge is either concave or convex but

cannot be convex along part of its length and concave elsewhere when there is no intervening edge junction.

The paper echoes his earlier paper in claiming that the interpretation of complex pictures requires "*a parsing operation on the results of context-free interpretation of picture fragments*" i.e. picture elements and their relationships need to be described, as a basis for interpreting the picture.

[This not a comprehensive summary of the contents of the 1971 paper.]

[... summary to be expanded ...]

The Acknowledgments section thanks R. Stanton, A. Sloman and especially Jack Lang, "for exposing deficiencies in the formulation by attempting to program earlier versions of the algorithm".

- **Controlled hallucination?**

Several AI vision publications refer to [the 1971 paper](#) as the source of the slogan "Perception is controlled hallucination", usually attributed to Max Clowes. However, That wording does not occur in this paper. The closest thing I have found is the wording in his 1969 paper quoted [above](#):

*"We may summarise the foregoing argument as 'People see what they expect to see'.
The essential rider is that what they want to see is things not pictorial relationships,*

That earlier slogan has more content insofar as it indicates explicitly that the hallucinated contents concern things depicted, not contents of the depiction.

NOTE on Hippolyte Taine

Keith Oatley informs me that a French author expressed a similar view of perception much earlier: Hippolyte Taine, a French discursive writer of the nineteenth century wrote:

"So our ordinary perception is an inward dream, which happens to correspond to things outside; and, instead of saying that a hallucination is a perception that is false, we must say that perception is a hallucination that is of the truth"

Taine, H. (1882). *De l'intelligence*, Tome 2. Paris: Hachette. (p. 13, emphasis in original). Full text, in English, available free here:
<http://hdl.handle.net/2027/uiuo.ark:/13960/t23b5zd2j>

Taine's formulation does not include the constructive role of perceptual mechanisms implicitly using prior knowledge of constraints on what is possible in the environment.

- **IJCAI 1971**

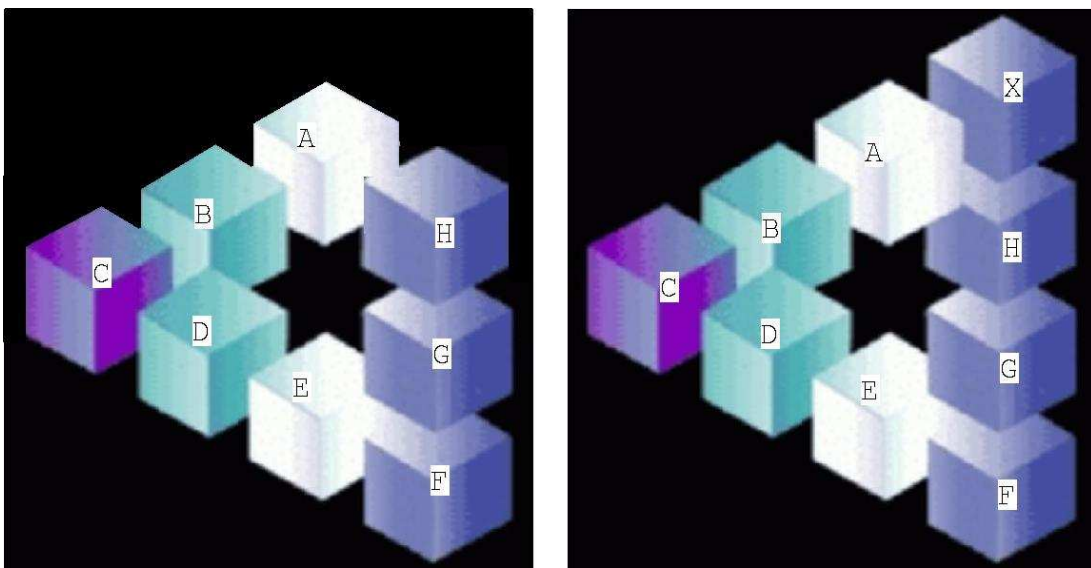
The second International Joint Conference on AI (IJCAI) was held at Imperial College in London in 1971. As a leading AI researcher in the host country, Max had an important organising role for that conference. One consequence was that he bullied me to fight a bout of 'flu before the submission deadline and submit a paper, which I did, and it changed my life.

The paper had an important relationship to Max's ideas. I had long been interested in the role of diagrams in mathematical reasoning, especially in Euclidean geometry, and also topology, logic and arithmetic, and my 1962 DPhil thesis was an attempt to defend Kant's

view that important kinds of mathematical reasoning could produce knowledge that was not empirical yet not just a matter of definitions and their logical consequences.

Max had observed that some of the details of human perception of pictures could be inferred from what we found ambiguous, synonymous, or anomalous, as explained [above](#) in connection with Chomsky's influence on his ideas. A visual "anomaly" occurs when 2D pictures have *parts* that are capable of representing 3D objects while the *whole* picture is incapable of doing so, just as some phrases or sentences have parts with normal semantic content whereas the whole phrase or sentence cannot, e.g. "John is (entirely) in the kitchen and (entirely) outside the kitchen". Pictorial examples include the Penrose triangle, the "devil's pitchfork" and various others. The connection between inferences and contradictions is well known in the case of sentences: e.g. if the joint truth of A and B is incompatible with the truth of C then A and B together imply not-C, and *vice versa*. This idea can be extended to contents of pictures.

I don't know if Max made that connection before we discussed it at Sussex University, though several AI researchers have, and have studied diagrammatic reasoning as an alternative to logical or algebraic deduction, e.g. as reported in [Glasgow et al. \(eds\) 1995](#), and elsewhere.



The image on the left depicts a possible 3-D scene.

Modifying it as on the right produces a picture that, if interpreted using the same semantic principles, represents an impossible 3-D scene (where blocks A, B, C form a horizontal line, blocks F, G, H form a vertical line, D and E are between and on the same level as C and F, and the new block X is co-linear with A, B, and C, and also with F, G, and H -- impossible!).

The drawing on the right was by Swedish artist, Oscar Reutersvard, in 1934

<http://im-possible.info/english/articles/triangle/triangle.html>

<http://butdoesitfloat.com/A-father-to-impossible-figures>

So a complex picture made of parts representing possible 3-D configurations may have N parts such that if a certain part X is added (e.g. an extra line joining two of the junctions, or a picture of an extra block that is simultaneously co-linear with two other linear groups, as in the above figure), then it becomes anomalous and cannot represent a 3-D configuration using the same rules of interpretation (based roughly on reversing projections from 3-D to 2-D). In other words the original N parts have a joint interpretation that entails that the situation depicted by adding the part X cannot exist. (I've left out some complications, including the fact that the depicted relationship

between blocks A and H has also changed. If the blocks were depicted spread out more in space, so that they are not overlapping, this change would not be necessary.)

XX

This is partly analogous to logical reasoning where N consistent propositions entail that an additional proposition X is false. So its negation can be inferred to be true. This example could not be handled by the Huffman-Clowes system, as it requires a richer grasp of geometry. Humans can reason that the configuration on the right is impossible without knowing any of the actual distances or sizes, whereas I don't believe any current AI vision system can do that. This is one of very many forms of geometrical and topological reasoning that are still beyond the scope of AI systems. Moreover I don't think neuroscientists have any idea how brains can support this kind of reasoning.

Note:

The point about impossible structures can be extended to impossible processes. An interesting rotating impossible triangle demonstration, by Hjalmar Snoep, can be found here. http://www.youtube.com/watch?v=wfdMRn_euXY

Further discussion of perception of impossibilities is available here, including a more detailed analysis of the Reutersvard picture:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html>

As far as I know, Max provided no explanation of how such impossibilities (anomalies) are recognized by humans. What cognitive mechanisms and processes underly the intuitions on which the linguistic and the geometric examples depend remains an unanswered question, though automated theorem provers can deal with logical and arithmetical inferences. The "line labelling" algorithm presented in Clowes (1971), which had been independently discovered by David Huffman, enabled a computer to derive the impossibility of certain 3-D interpretations of picture fragments from the fact that implied combinations of edge and junction features that were ruled out by programmer supplied rules. Those rules were discovered by humans thinking about 3D volumes bounded by planar surfaces. As yet I don't know of any AI program that can discover such geometrical constraints itself, though there may be some mathematically equivalent arithmetical theorem about coordinates of collinear or co-planar sets of points that a machine could prove. Note that training a program on a huge collection of pictures in which the anomalous cases do not occur might give the program the ability to assign a *low probability* to such cases, but would not give it a human-like understanding that they are *impossible*, and why they are.

There have been extensions of the Huffman-Clowes idea to reasoning about curved objects, and scenes involving more complex polyhedra along with cracks and shadows, e.g. Waltz.

Moreover the ideas were extended by Waltz, Hinton, Mackworth and others to allow constraint-propagation techniques to be used to improve efficiency, and to allow use of 'soft' constraints so that missing or extra features caused by poor lighting or use of low quality cameras could allow most of the image to be interpreted correctly despite missing or spurious fragments. In all these cases human reasoning was used to discover the rules to be used, rather than reasoning by machines with understanding of spatial structures.

My hope in 1971 was that we would one day understand how to build a 'baby' robot with abilities to interact with its environment and with a suitable set of innate mechanisms for extending its knowledge and modes of reasoning in something like the way human children and mathematicians do. This could explain why Kant was right about the nature of

mathematical knowledge, by demonstrating a working system, able to use spatial reasoning to make a variety of deep mathematical discoveries, including a significant subset of topology, Euclidean geometry and arithmetic. I also thought that in some contexts such non-logical forms of reasoning would not only be closer to human mathematical reasoning than logical deductions from axioms, but might add additional power to intelligent machines, in particular heuristic power based on search spaces closer to the problem domain. I don't know whether Max agreed with this, but it was he who pressed me in 1971 to submit the paper criticising purely logicist AI that was accepted for IJCAI.

Since writing the paper I have discovered that it is very difficult to get computers to emulate the forms of spatial reasoning that are common in human mathematics, although some special cases can be handled by translating geometry into arithmetic and algebra following Descartes and using modern logic. But that's not the method used by Euclid centuries before Descartes, Frege and modern logicians. (For anyone interested, some topological examples are discussed [here](#).) There are also no working models capable of explaining or replicating the spatial reasoning abilities of nest building birds (e.g. weaver birds), elephants, human toddlers, and many other intelligent animals.

- **M.B. Clowes, "Scene analysis and picture grammars", (1972)**

in: F. Nake and A. Rosenfeld, eds., *Graphic Languages* (Proc. IFIP Working Conf on Graphic Languages, Vancouver, 1972), Amsterdam 1972, pp. 70-82; also in: *Machine Perception of Patterns and Pictures*, Institute of Physics, London and Bristol, 1972, pp. 243-256

(I have not found an online version of this.)

- **Move from EP to the Arts and Social Studies (1973)**

Partly as a result of Max's dislike of the proximity (and smell) of rats, along with his interest in linguistics, and psychology of humans, he was offered a chair in AI, in the Arts and Social studies area, of the university, in 1973, to help start the new interdisciplinary Cognitive Studies Programme, mentioned in the obituary notice. This was located in the School of Social Sciences (SOCS).

- **M.B. Clowes, ICA Lecture and book chapter 'Man the creative machine: A perspective from Artificial Intelligence research'(1972/3)**

in: J. Benthall (ed) *The Limits of Human Nature*, Allen Lane, London, 1973.

NOTE added 26 Feb 2016:

This section has been slightly expanded, with additional quotations from the paper.

A talk with this title was given at the Institute of Contemporary Arts (ICA), as part of a series of lectures at the ICA in 1971-2. Papers corresponding to the lectures were collected in the book edited by Jonathan Benthall (which also included a paper on AI language processing by Terry Winograd). Although the phrase "controlled hallucination" is not in Max's paper, the idea is there. This paper also emphasises the role of metaphor in perception:


"The sense of a picture is that web of knowledge-based relations between objects and situations, including the viewer, which is maximally supported by the picture." (p.204)

and later on that page

"The point I am making here has been made more clearly and at greater length by the art historian Ernst Gombrich in his book Art and Illusion. It is, briefly, that all comprehension is essentially metaphorical; what we mean by 'literal' is merely that the sense is consistent with the total context of the act of perception and not unrelated to, or at variance with, it."

He emphasises that this would be impossible without "heterarchical" processing, i.e. a mixture of bottom-up, top-down and middle-out processing. (Later demonstrated, partly under Max's influence in the Popeye program described in [Chapter 9 of my 1978 book](#) e.g. [Figure 9.6](#))

The ICA talk included a still from John Schlesinger's film "Sunday bloody Sunday" in which two intertwined bodies were shown, presenting the viewer with the task of using world knowledge to decide which parts belonged to the same body: an example of creative "controlled hallucination" of hidden connections. The same picture is in the published paper. My amateurish sketch of the scene is below.



Which hands belong to whom?

Answering this question requires use of knowledge of human anatomy to hallucinate unseen connections.

However I acquired the book in April 2014, and searched through the paper. There is no mention of hallucination, though Max may have used the phrase "controlled hallucination" when presenting the talk at ICA. Readers should be able to use their own controlled creativity to hallucinate invisible connections in order to work out which hands belong to which person, using knowledge of human anatomy.

- **O’Gorman and Clowes Collinearity (1973, 1976)**

Frank O’Gorman, M. B. Clowes: Finding Picture Edges through Collinearity of Feature Points. IJCAI 1973: 543-555

That conference paper was later revised for a journal:

Frank O’Gorman, M. B. Clowes: Finding Picture Edges Through Collinearity of Feature Points. IEEE Trans. Computers 25(4): 449-456 (1976)

https://www.researchgate.net/publication/3046601_Finding_Picture_Edges_Through_Collinearity_of_Feature_Points

- **Other publications to be added:** Inaugural Lecture (197??) (unpublished)

Frank O’Gorman: Edge Detection Using Walsh Functions. *Artif. Intell.* 10(2): 215-223 (1978)

- **Students and colleagues at Sussex**

Max had several other research students at Sussex working on various aspects of vision, and he also supervised some AI MSc projects. One of his students is [Alan Mackworth](#) who went on to become a leading AI researcher in Canada with a prominent international role, e.g. in IJCAI and AAAI, among others.

[Steve Draper](#) followed up a BSc in Physics and a DPhil in AI at Sussex supervised by Max, with a career in psychology. A paper based on his DPhil work (funded by one of Max’s projects) was published in 1980:

<http://www.cs.bham.ac.uk/research/projects/cogaff/aisb1980/draper-aisb1980.pdf>

S. Draper, Using Models to Augment Rule-Based Programs, in *Proceedings of The AISB-80 Conference on Artificial Intelligence* Amsterdam 4-8 July, 1980, available here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/aisb1980/aisb1980.pdf>

Larry Paul was supervised by both Max and Aaron Sloman.

Frank O’Gorman, after an MSc in Computer Science at Birmingham, worked with Max as a research assistant for several years (and also taught me a great deal about programming and computer science, partly by teaching me Algol 68, which was used for their research in the mid 1970s). After the grant ended he worked for a while on my POPEYE project described [here](#), which attempted to extend some of Max’s ideas. The grant was awarded by the Cognitive Science panel of the Science Research Council on condition that Max had an advisory role. Others on the team were David Owen, and, for a short time, Geoffrey Hinton. We were all deeply influenced by Max. [Other students, collaborators, etc. ... ?]

- **Obituary notices**

- The Sussex University Bulletin obituary notice in 1981:
<http://www.sussex.ac.uk/internal/bulletin/downloads/1980-1989/1981/May/19810512.pdf> mentions that Max was at Reading, Oxford, and NPL, before he went to Australia.
- There was a short notice in the proceedings of IJCAI 1981, for which I do not have the text.
- Experiencing Computation: A Tribute to Max Clowes, by Aaron Sloman, was published in *Computing in Schools* in 1981, the same year as the BBC Micro was announced.
- The longer version included above was published in 1984, in *New horizons in educational computing*, Ed. Masoud Yazdani, Ellis Horwood Series In Artificial Intelligence, pp. 207--219, Chichester,
<http://www.cs.bham.ac.uk/research/projects/cogaff/00-02.html#71>

REFERENCES

- M.L.J. Abercrombie *The Anatomy of Judgement* (1960), New York: Basic Books, 1960. Summarised [here](#),

- Margaret A. Boden, Review of *Machine Intelligence 4* Eds. Bernard Meltzer, Donald Michie in *The British Journal for the Philosophy of Science*, Vol. 21, No. 2 (May, 1970), pp. 212-214 (OUP, Brit.Soc.Phil.Sci) <http://www.jstor.org/stable/686707>
 - <http://dx.doi.org/10.1007/s13347-013-0115-x> Margaret A. Boden, "Remarks on Receiving the Covey Award", in *Philosophy & Technology*, 26, 3, 2013, Springer, Netherlands, pp333-339,
 - David A. Huffman, "Impossible objects as nonsense sentences", in *Machine Intelligence 6* Meltzer, B., and Michie, D. (Eds.), Edinburgh University Press, Edinburgh (1971), 295-323.
 - Eds. J. Glasgow and H. Narayanan and B. Chandrasekaran, *Diagrammatic Reasoning: Computational and Cognitive Perspectives*, MIT Press, Cambridge, MA, 1995,
-

About this document ...

This document was originally generated, using **LaTeX2HTML**, by Aaron Sloman on 2001-02-17 followed by some hand editing. Considerable further editing is indicated by notes at the top and in-line comments.

Installed: 17 Feb 2001

Updated: Formatting changes. Major and minor textual changes.

29 Jul 2016: Minor formatting changes.

26 Feb 2016: Added more quotes from ICA paper.

9 Feb 2016: updated address for blocks figure.

31 May 2015: slightly modified Reutersvard figure

1 May 2014; 5 May 2014; 11 May 2014; 7 Sep 2014; 12 Sep 2014